To date the Internet has done fine with respect to providing best-effort connectivity, but it's managed to do so specifically *because* it has been largely protocol- and user-agnostic. This inherent neutrality has allowed it to grow wildly, far beyond what its inventors ever envisioned. While I believe we still have a ways to go before we reach the prophesied "convergence" of devices (computers, phones, entertainment devices all merged into a single device which performs all tasks equally well), I think we're clearly at a convergence for data transport. Ensuring that the infrastructure to move this data exists, and does so *without* discrimination determined by private entities, is extremely important.

In order to maintain this flexibility, the conduits for data should be at *LEAST* as open and non-discriminatory as the conduits for voice traditionally have been. Filtering based on content in the digital realm strikes me as analogous to a telco that provides "basic phone connectivity" that only passes frequencies in the range of 140-180Hz, then charging more for "full dynamic range telephony." It is with these statements in mind that I would like to try to address the Commission's questions regarding appropriate packet/payload prioritization.

There are few examples better than the University of Minnesota's core network infrastructure for illustrating the difficulty in prioritizing network traffic. UMN's network is a microcosm in itself of the Internet, and contains a widely varied mix of traffic. Much of it is very common, e.g., WWW, e-mail, instant messaging, etc. Some of it is less common, e.g., IPSEC or similar traffic. Within the core network, little if any prioritization is done; traffic is passed unhindered. Limiting that's applied is done as close to the edge as possible. For example, traffic shaping is applied at the reshall border to prevent those hosts. Connectivity is still fast (I believe the dorms may be wired for gigabit), but outbound filtering is generally very weak; those hosts can connect to whatever outside server they choose. (Inbound services to dorm room machines is prohibited by policy.)

On the other hand, connectivity to UMN-owned hosts is largely uncontrolled. Filtering is still edge-driven, but there's very little filtering or packet shaping. This is largely because it is impossible to accurately predict what protocols or connectivity will be needed for

research or other purposes. Some packet shaping is applied, and heavy traffic generators are sometimes put into their own carved-out sections of network, but the majority of traffic is passed unfiltered. This is as it should be.

Similarly, an ISP has (or should have) a responsibility to ensure that its customers' data is passed. If the customer and ISP agree that the ISP is going to limit connectivity for some reason (e.g., for some security purpose), that's OK, but the default stance should be to *allow* connectivity. ISPs, like the central networking people at a large university, are in an extremely poor place to determine what should be filtered and why. The usual excuses offered by ISPs for doing so are capacity and expense, security, and some variant on "think of the children."

Capacity and planning is clearly an issue. An ISPs equipment can only carry a finite amount of data in a given time period. However, I think these issues are less due to a need for disregarding neutral data transport policies, and more about a) a failure to do proper capacity planning and trend analysis; and b) overselling/oversubscribing existing infrastructure. If an ISP says they'll provide a 2Mb/sec pipe, then they should be able to deliver.

Security is another specious argument used in trying to sell the idea of network discrimination. On the one hand, it's pretty clear that a lot of people have zero ability to secure their computers. (Consider the FBI's recent "Operation 'Bot Roast" announcement.) On the other hand, an ISP is in the exact same situation as UMN's central networking group: They can't *possibly* determine accurately and quickly what is or is not "safe." New protocols, and new uses for existing protocols, crop up all the time, and it's a target that moves far too fast for an ISP to track.

The last common argument is related to the security argument, but uses children as the reason. I think the Commission will be in full agreement with me that child pornography is a terrible, hideous crime, and that the perpetrators of such crime should receive the full punishment of the law. That said, there is no feasible way an ISP can automatically identify and block child-unfriendly traffic on an ongoing basis. It's a problem with an impossible solution given today's

technology and legal environment.

I've tried to keep these comments brief, but hope that they've still been of value to the Commission. I appreciate the opportunity to comment, and welcome any inquiries the Commission may have on this matter.

Thanks for your time.

--
Alan Amesbury
amesbury@umn.edu